

IN THE CLAIMS:

Please write the claims to read as follows:

1 1. (Currently Amended) A method for a storage operating system implemented in a
2 storage system to optimize an amount of readahead data retrieved from
3 a data container of the storage system, the method comprising:
4 maintaining a plurality of files;
5 maintaining, for a selected file of the plurality of files, a plurality of readset data
6 structures, each readset data structure holding a plurality of factors for a selected
7 readstream, the plurality of factors allowing the system to adjust adaptively the amount of
8 data retrieved from the data container;
9 receiving a client read request for a particular read stream at the storage system;
10 locating a readset data structure for the particular read stream;
11 determining whether the storage operating system is permitted to retrieve
12 readahead data from the data container in response to the received client read request; and
13 if it is determined that the storage operating system is permitted to retrieve
14 readahead data from the data container, performing the steps of,
15 (i) ~~selecting~~ adjusting an amount of readahead data to retrieve from the data
16 container, based on the plurality of factors stored within the readset data structure;
17 (ii) retrieving the adjusted ~~selected~~ amount of readahead data from the data
18 container; and
19 (iii) ~~determining if the readahead-readset data structure meets a criteria for being~~
20 ~~updated, and if the readahead-readset data structure meets the criteria then updating the~~
21 ~~readahead-readset data structure.~~

1 2. (Original) The method of claim 1, wherein the data container is a file, directory, vdisk
2 or lun.

1 3. (Original) The method of claim 1, wherein the storage operating system is determined
2 to be permitted to retrieve readahead data from the data container when the client-
3 requested data extends the read stream past a predetermined next readahead value.

1 4. (Original) The method of claim 3, wherein the predetermined next readahead value is
2 stored in a readset data structure associated with the read stream.

1 5. (Currently Amended) The method of claim 3, wherein the predetermined next
2 readahead value is updated based on a percentage of the ~~selected~~ adjusted amount of
3 readahead data.

1 6. (Previously Presented) The method of claim 1, wherein a read-access style associated
2 with the data container is one of the plurality of factors used to select the amount of
3 readahead data.

1 7. (Currently Amended) The method of claim 6, wherein the ~~selected~~ adjusted amount
2 of readahead data equals zero if the read-access style corresponds to a random read-
3 access style.

1 8. (Previously Presented) The method of claim 1, wherein a number of client read
2 requests processed in the read stream is one of the plurality of factors used to select the
3 amount of readahead data.

- 1 9. (Original) The method of claim 8, wherein the number of client read requests
2 processed in the read stream is stored as a count value in a readset data structure
3 associated with the read stream.
- 1 10. (Previously Presented) The method of claim 1, wherein the amount of client-
2 requested data is one of the plurality of factors used to select the amount of readahead
3 data.
- 1 | 11. (Currently Amended) The method of claim 10, wherein the ~~selected~~-adjusted amount
2 of readahead data is set equal to a predetermined upper limit for large amounts of client-
3 requested data.
- 1 | 12. (Currently Amended) The method of claim 1, wherein the ~~selected~~-adjusted amount
2 of readahead data is doubled if the number of client read requests processed in the read
3 stream is greater than a first threshold value.
- 1 13. (Original) The method of claim 1, wherein the client-requested data is identified as
2 read-once data when either (i) the number of client read requests processed in the read
3 stream is greater than a second threshold value or (ii) a set of metadata associated with
4 the read stream indicates that the client-requested data is read-once data.
- 1 | 14. (Currently Amended) The method of claim 1, wherein the ~~selected amount~~adjusted
2 amount of readahead data is stored in one or more buffers enqueued on a flush queue, the
3 flush queue being configured to reuse buffers after a predetermined period of time.
- 1 15. (Original) The method of claim 14, wherein the predetermined period of time equals
2 two seconds.

1 16. (Currently Amended) An apparatus configured to implement a storage operating
2 system that optimizes an amount of readahead data retrieved from a data container of the
3 apparatus, the apparatus comprising:

4 means for maintaining a plurality of files;

5 means for maintaining, for a selected file of the plurality of files, a plurality of
6 readset data structures, each readset data structure holding a plurality of factors for a
7 selected readstream, the plurality of factors allowing the system to adjust adaptively the
8 amount of data retrieved from the data container;

9 means for receiving a client read request for a particular read stream;

10 means for locating a readset data structure for the particular read stream;

11 means for determining whether the storage operating system is permitted to
12 retrieve readahead data from the data container in response to the received client read
13 request;

14 means for ~~selecting and~~ adjusting an amount of readahead data to retrieve from the
15 data container based on the plurality of factors stored within the readset data structure;

16 ~~and~~

17 means for retrieving the ~~selected amount~~ adjusted amount of readahead data from
18 the data container; and

19 means for determining if the readahead data structure meets a criteria for being
20 updated, and if the readahead data structure meets the criteria then updating the
21 readahead data structure.

1 17. (Original) The apparatus of claim 16, wherein the data container is a file, directory,
2 vdisk or lun.

1 18. (Original) The apparatus of claim 16, wherein the storage operating system is
2 determined to be permitted to retrieve readahead data from the data container when the
3 client-requested data extends the read stream past a predetermined next readahead value.

1 19. (Currently Amended) The apparatus of claim 18, further comprising means for
2 updating the predetermined next readahead value based on a percentage of the ~~selected~~
3 amount~~adjusted amount~~ of readahead data.

1 20. (Previously Presented) The apparatus of claim 16, wherein the plurality of factors
2 used to select the amount of readahead data includes at least one of:

- 3 (i) the amount of client-requested data,
4 (ii) a number of client read requests processed in the read stream, and
5 (iii) a read-access style associated with the data container.

1 21. (Currently Amended) The apparatus of claim 16, wherein the ~~selected~~
2 amount~~adjusted amount~~ of readahead data is doubled if the number of client read requests
3 processed in the read stream is greater than a first threshold value.

1 22. (Currently Amended) A storage system configured to optimize an amount of
2 readahead data retrieved from a data container of the storage system, the storage system
3 comprising:

- 4 a network adapter for receiving a client read request for a particular read stream;
5 a memory configured to store instructions for implementing a storage operating
6 system that performs the steps of:
7 locating a readset data structure for the particular read stream;

8 ~~selecting~~adjusting, in response to the readset data structure for the particular read
9 stream, ~~a the selected amount~~adjusted amount of readahead data from the data container
10 based on a plurality of factors, the plurality of factors allowing the system to adjust
11 adaptively the amount of data retrieved from the data container;
12 retrieving the ~~selected amount~~ adaptively adjusted amount of readahead data from
13 the data container.

1 23. (Original) The storage system of claim 22, wherein the data container is a file,
2 directory, vdisk or lun.

1 24. (Original) The storage system of claim 22, wherein the storage operating system is
2 determined to be permitted to retrieve readahead data from the data container when the
3 client-requested data extends the read stream past a predetermined next readahead value.

1 25. (Currently Amended) The storage system of claim 24, wherein the predetermined
2 next readahead value is updated based on a percentage of the ~~selected amount~~adjusted
3 amount of readahead data.

1 26. (Previously Presented) The storage system of claim 22, wherein the plurality of
2 factors used to select the amount of readahead data includes at least one of:

- 3 (i) the amount of client-requested data,
4 (ii) a number of client read requests processed in the read stream, and
5 (iii) a read-access style associated with the data container.

27. (Currently Amended) The storage system of claim 22, wherein the ~~selected~~
~~amount~~adjusted amount of readahead data is doubled if the number of client read requests
processed in the read stream is greater than a first threshold value.

28. (Currently Amended) A computer-readable media comprising instructions for
execution in a processor for the practice of a method for a storage operating system
implemented in a storage system to optimize an amount of readahead data retrieved from
a data container of the storage system, the method comprising:

maintaining a plurality of files;

maintaining, for a selected file of the plurality of files, a plurality of readset data
structures, each readset data structure holding a plurality of factors for a selected
readstream based on a plurality of factors, the plurality of factors allowing the system to
adjust adaptively the amount of data retrieved from the data container;

receiving a client read request for a particular read stream at the storage system;

locating a readset data structure for the particular read stream;

determining whether the storage operating system is permitted to retrieve
readahead data from the data container in response to the received client read request;

if it is determined that the storage operating system is permitted to retrieve
readahead data from the data container, performing the steps of:

(i) ~~selecting and~~adjusting an amount of readahead data to retrieve from the data
container based on the plurality of factors stored within the readset data structure; and

(ii) retrieving the ~~selected amount~~adjusted amount of readahead data from the
data container;

(iii) determining if the ~~readahead-readset~~ data structure meets a criteria for being
updated, and if the ~~readahead-readset~~ data structure meets the criteria then updating the
~~readahead-readset~~ data structure .

1 29. (Original) The computer-readable media of claim 28, wherein the data container is a
2 file, directory, vdisk or lun.

1 30. (Previously Presented) The method of claim 1, wherein the retrieved readahead data
2 is stored in one or more buffers, the buffers containing a flush queue, the flush queue
3 being configured to reuse buffers after a predetermined period of time.

1 31. (Previously Presented) The method of claim 30, wherein the read stream
2 corresponds to a read-once data transfer and data retrieved from the data container is
3 stored in the flush queue.

1 32. (Previously Presented) The method of claim 30, wherein the retrieved readahead
2 data is stored in the flush queue.

1 33. (Previously Presented) The method of claim 30, wherein one or more buffers
2 accessed from the flush queue are re-enqueued on a normal queue.

1 34. (Currently Amended) A method for optimizing readahead data retrieved from
2 a data container of a storage system, the method comprising:
3 maintaining a plurality of files;
4 maintaining, for a selected file of the plurality of files, a plurality of readset data
5 structures, each readset data structure holding a plurality of factors for a selected
6 readstream;
7 receiving a client read request for a particular read stream;
8 locating a readset data structure for the particular read stream;

9 | ~~selecting and adjusting an~~ amount of readahead data based on the indicated amount
10 | of client-requested data and a corresponding readset data structure; and
11 | ~~retrieving the selected amount~~ adjusted amount of readahead data from the data
12 | container.

1 | 35. (Currently Amended) The method of claim 34, wherein the ~~selected amount~~ adjusted
2 | amount of readahead data is set equal to a multiple of a predetermined amount, and
3 | wherein the multiple is associated with the amount of client-requested data.

1 | 36. (Currently Amended) The method of claim 34, wherein the ~~selected amount~~ adjusted
2 | amount of readahead data is set equal to a multiple of the amount of client-requested data.

1 | 37. (Currently Amended) The method of claim 36, further comprising the step of
2 | rounding the ~~selected amount~~ adjusted amount of readahead data to the size of a data
3 | block.

1 | 38. (Currently Amended) The method of claim 34, wherein the ~~selected amount~~ adjusted
2 | amount of readahead data is set equal to a predetermined upper limit.

1 | 39. (Currently Amended) A method for optimizing readahead data retrieved from a data
2 | container of a storage system, the method comprising:

3 | maintaining a plurality of files;

4 | maintaining, for a selected file of the plurality of files, a plurality of readset data
5 | structures, each readset data structure holding a plurality of factors for a selected
6 | readstream based on a plurality of factors, the plurality of factors allowing the system to
7 | adjust adaptively the amount of data retrieved from the data container;

8 receiving a client read request for a particular read stream at the storage system;
9 locating a readset data structure for the particular read stream;
10 ~~selecting for~~adjusting for the read stream an amount of readahead data based on a
11 read-access style associated with the data container and a corresponding readset data
12 structure; and
13 ~~retrieving the selected amount~~adjusted amount of readahead data from the data
14 container.

1 40. (Currently Amended) The method of claim 39, wherein the ~~selected amount~~adjusted
2 amount of readahead data equals zero if the read-access style corresponds to a random
3 read-access style.

1 41. (Currently Amended) A method for optimizing readahead data retrieval from
2 a data container of a storage system associated with a number of storage devices, the
3 method comprising:
4 maintaining a plurality of files;
5 maintaining, for a selected file of the plurality of files, a plurality of readset data
6 structures, each readset data structure holding a plurality of factors for a selected
7 readstream based on a plurality of factors, the plurality of factors allowing the system to
8 adjust adaptively the amount of data retrieved from the data container;
9 receiving a client read request for a particular read stream at the storage system;
10 locating a readset data structure for the particular read stream;
11 ~~selecting~~adjusting an amount of readahead data based on a number of storage
12 devices and a corresponding readset data structure; and
13 ~~retrieving the selected amount~~adjusted amount of readahead data from the data
14 container.

1 42. (Currently Amended) The method of claim 41, wherein ~~the step of selecting~~
2 ~~an~~adjusting an amount of readahead data further comprises:

3 determining whether a flag is associated with the read stream, the flag indicating
4 that the storage system is associated with more than a predetermined number of storage
5 devices; and

6 in response to determining whether the flag is associated, ~~selecting the~~adjusting
7 the amount of readahead data.

1 43. (Previously Presented) The method of claim 41, wherein the storage devices
2 comprise one or more disks.

1 44. (Currently Amended) A method for optimizing readahead data retrieval in a storage
2 system, the method comprising:

3 maintaining a plurality of files;

4 maintaining, for a selected file of the plurality of files, a plurality of readset data
5 structures, each readset data structure holding a plurality of factors for a selected
6 readstream based on a plurality of factors, the plurality of factors allowing the system to
7 adjust adaptively the amount of data retrieved from the data container;

8 receiving a client read request for a particular read stream at the storage system;

9 locating a readset data structure for the particular read stream;

10 ~~selecting an~~adjusting an amount of readahead data based on a plurality of factors
11 stored within a corresponding readset data structure; and

12 retrieving the ~~selected amount~~adjusted amount of readahead data from a data
13 container.

1 45. (Previously Presented) The method of claim 44, wherein the retrieved readahead
2 data is stored in one or more buffers, the buffers containing a flush queue, the flush queue
3 being configured to reuse buffers after a predetermined period of time.

1 46. (Previously Presented) The method of claim 45, wherein the read stream
2 corresponds to a read-once data transfer and data retrieved from the data container is
3 stored in the flush queue.

1 47. (Previously Presented) The method of claim 45, wherein the retrieved readahead
2 data is stored in the flush queue.

1 48. (Previously Presented) The method of claim 45, wherein one or more buffers
2 accessed from the flush queue are re-enqueued on a normal queue.

1 49. (Currently Amended) A system for optimizing readahead data retrieval in a storage
2 system, the system comprising:

3 maintaining a plurality of files;

4 maintaining, for a selected file of the plurality of files, a plurality of readset data
5 structures, each readset data structure holding a plurality of factors for a selected
6 readstream based on a plurality of factors, the plurality of factors allowing the system to
7 adjust adaptively the amount of data retrieved from the data container;

8 means for receiving a client read request for a particular read stream at the storage
9 system;

10 means for locating a readset data structure for the particular read stream;

11 means for selecting and adjusting an amount of readahead data based on a plurality
12 of factors stored within a corresponding readset data structure; and

13 | means for retrieving the ~~selected amount~~adjusted amount of readahead data from
14 | a data container.

1 50. (Previously Presented) The system of claim 49, wherein the retrieved readahead data
2 is stored in one or more buffers, the buffers containing a flush queue, the flush queue
3 being configured to reuse buffers after a predetermined period of time.

1 51. (Previously Presented) The system of claim 50, wherein the read stream corresponds
2 to a read-once data transfer and data retrieved from the data container is stored in the
3 flush queue.

1 52. (Previously Presented) The system of claim 50, wherein the retrieved readahead data
2 is stored in the flush queue.

1 53. (Previously Presented) The system of claim 50, wherein one or more buffers
2 accessed from the flush queue are re-enqueued on a normal queue.

1 54. (Currently Amended) A method, comprising:
2 maintaining a plurality of files;
3 maintaining, for a selected file of the plurality of files, a plurality of readset data
4 structures, each readset data structure holding a plurality of factors for a selected
5 readstream based on a plurality of factors, the plurality of factors allowing the system to
6 adjust adaptively the amount of data retrieved from the data container;
7 receiving a plurality of client read requests for a particular read stream at a storage
8 system;
9 locating a readset data structure for the particular read stream;

10 | ~~selecting and adjusting an~~ amount of readahead data to retrieve from one or more
11 | data containers based on a plurality of factors stored within the readset data structure;
12 | retrieving the ~~selected amount~~ adjusted amount of readahead data from the data
13 | container;
14 | processing one or more of the plurality of client read requests; and
15 | ~~selecting and adjusting~~, as client requests are processed, the plurality of factors stored
16 | within the readset data structure associated with each read stream to optimize the amount
17 | of readahead data that is cached for each read stream.

1 55. (Previously Presented) The method of claim 54, further comprising:
2 determining whether the storage operating system is permitted to retrieve
3 readahead data from the one or more data containers in response to each received client
4 read request.

1 56. (Previously Presented) The method of claim 54, wherein the one or more data
2 containers are at least one of a file, a directory, a vdisk or a lun.

1 57. (Previously Presented) The method of claim 55, wherein the storage operating
2 system is determined to be permitted to retrieve readahead data from the one or more data
3 containers when the client-requested data extends the read stream past a predetermined
4 next readahead value.

1 58. (Previously Presented) The method of claim 57, wherein the predetermined next
2 readahead value is stored in a readset data structure associated with the read stream.

1 59. (Currently Amended) The method of claim 57, wherein the predetermined next
2 readahead value is updated based on a percentage of the ~~selected amount~~adjusted amount
3 of readahead data.

1 60 (Previously Presented) The method of claim 54, wherein a read-access style
2 associated with the one or more data containers is one of the plurality of factors used to
3 select the amount of readahead data.

1 61. (Currently Amended) The method of claim 60, wherein the ~~selected amount~~adjusted
2 amount of readahead data equals zero if the read-access style corresponds to a random
3 read-access style.

1 62. (Previously Presented) The method of claim 54, wherein a number of client read
2 requests processed in the read stream is one of the plurality of factors used to select the
3 amount of readahead data.

1 63. (Previously Presented) The method of claim 62, wherein the number of client read
2 requests processed in the read stream is stored as a count value in a readset data structure
3 associated with the read stream.

1 64. (Previously Presented) The method of claim 54, wherein the amount of client-
2 requested data is one of the plurality of factors used to select the amount of readahead
3 data.

1 65. (Currently Amended) The method of claim 64, wherein the ~~selected amount~~adjusted
2 amount of readahead data is set equal to a predetermined upper limit for large amounts of
3 client-requested data.

1 66. (Currently Amended) The method of claim 54, wherein the ~~selected amount~~adjusted
2 amount of readahead data is doubled if the number of client read requests processed in
3 the read stream is greater than a first threshold value.

1 67. (Previously Presented) The method of claim 55, wherein the client-requested data is
2 identified as read-once data when either (i) the number of client read requests processed
3 in the read stream is greater than a second threshold value or (ii) a set of metadata
4 associated with the read stream indicates that the client-requested data is read-once data.

1 68. (Currently Amended) The method of claim 54, wherein the ~~selected amount~~adjusted
2 amount of readahead data is stored in one or more buffers enqueued on a flush queue, the
3 flush queue being configured to reuse buffers after a predetermined period of time.

1 69. (Currently Amended) A method for optimizing readahead data retrieval for a storage
2 system, the method comprising:

3 maintaining a plurality of files;

4 maintaining, for a selected file of the plurality of files, a plurality of readset data
5 structures, each readset data structure holding a plurality of factors for a selected
6 readstream;

7 receiving a client read request for a particular read stream at the storage system;

8 locating a readset data structure for the particular read stream;

9 ~~selecting an~~adjusting an amount of readahead data in response to a corresponding
10 readset data structure based on a plurality of factors, the plurality of factors allowing the
11 system to adjust adaptively the amount of data retrieved from the data container; and

12 | retrieving the ~~selected amount~~adaptively adjusted amount of readahead data from
13 | a data container.

1 | 70. (Previously Presented) The method of claim 69, further comprising:
2 | allocating more readsets for the file in response to processing one or more client
3 | “write” requests to the file.

1 | 71. (Currently Amended) A storage system, comprising:
2 | an operating system to maintain a plurality of files, and to maintain, for a selected
3 | file of the plurality of files, a plurality of readset data structures, each readset data
4 | structure holding a plurality of factors for a selected readstream;
5 | a network adapter to receive a client read request at the storage system for client-
6 | requested data stored in a file;
7 | a plurality of readset data structures associated with the client-requested data;
8 | an operating system to establish a read stream corresponding to each readset data
9 | structure of the plurality of readset data structures;
10 | a process to ~~select~~adjust an amount of readahead data in response to a
11 | corresponding readset data structure of the plurality of readset data structures which is
12 | based on a plurality of factors, the plurality of factors allowing the system to adjust
13 | adaptively the amount of data retrieved from the data container;
14 | a storage adapter to retrieve the ~~selected amount~~adjusted amount of readahead
15 | data from the data container; and
16 | the operating system to determine if the readset data structure meets a criteria for
17 | being updated, and if the ~~readahead-readset~~ data structure meets the criteria, then
18 | updating the ~~readahead-readset~~ data structure.

- 1 72. (Previously Presented) The storage system of claim 71, further comprising:
2 the operating system to allocate more readsets for the file in response to
3 processing one or more client “write” requests to the file.